

Explainable AI Framework for Health Monitoring and Fault Diagnosis in Safety- Critical Time-Series Systems

B. Ashwin Kumar, N. Thomas Peter, E. Benitha Sowmiya

U.G. Student, Department of Computer Science and Engineering, Bharath Institute of Science and Technology, Bharath
Institute of Higher Education and Research, Tambaram, Chennai, Tamil Nadu, India

U.G. Student, Department of Computer Science and Engineering, Bharath Institute of Science and Technology, Bharath
Institute of Higher Education and Research, Tambaram, Chennai, Tamil Nadu, India

Assistant professor, Department of Computer Science and Engineering, Bharath Institute of Science and Technology,
Bharath Institute of Higher Education and Research, Tambaram, Chennai, Tamil Nadu, India

ABSTRACT: This paper introduces an intelligent explainable AI framework designed for health monitoring and fault diagnosis in safety-critical systems that rely on time-series data. The system processes multivariate sensor readings to detect anomalies, trace their sources, and generate a dynamic health index that shows overall system condition. Traditional monitoring tools only raise alerts when problems appear, but this framework combines unsupervised machine learning with explainable AI methods to deliver clear, interpretable explanations of system behaviour. It handles data preprocessing, sliding-window segmentation, feature engineering, and anomaly detection through Isolation Forest and autoencoder models. A dedicated fault diagnosis step separates sensor-level issues from broader system-level problems. The health index, scored from 0 to 100, uses weighted sensor importance to track degradation over time. SHAP values provide transparent explanations of model decisions and highlight the most influential features

KEYWORDS: Explainable AI, Time-Series Analysis, Anomaly Detection, Health Monitoring, Fault Diagnosis, SHAP, Isolation Forest.

I. INTRODUCTION

Safety-critical systems, such as industrial control platforms in factories, telemetry networks in remote monitoring stations, or even large-scale IoT setups that track everything from temperature to voltage, constantly generate huge streams of multivariate time-series data. Every second, sensors send readings that form long sequences showing how the system is behaving over time. Keeping everything reliable and stopping failures before they happen requires nonstop analysis of this data. Yet reviewing all those numbers manually is slow, tiring, and full of human mistakes that can easily be missed during night shifts or busy periods.

Most current monitoring setups only use basic threshold alerts that ring a bell when a value crosses a fixed line, or they rely on hidden black-box machine learning models that flag problems without ever explaining the reason behind them. This leaves engineers in the dark, forcing them to spend extra time investigating instead of fixing the root issue. The framework in this paper changes that completely. It not only catches anomalies fast and accurately but also clearly explains why they appeared, while keeping a continuous, easy-to-read eye on overall system health. Built as a clean, modular package with separate parts that can be swapped or improved independently, it works with any time-series data in demanding real-life situations. The goal is simple: make advanced AI monitoring feel approachable, reliable, and actually helpful for the people who run these critical systems every day.

By focusing on both detection and understanding, the system bridges the gap between raw data and real-world decisions. It turns overwhelming sensor streams into clear stories that support faster troubleshooting, better maintenance planning, and higher overall safety

Volume 15, Issue 4, April 2026**|DOI: 10.15680/IJRSET.2026.1504227|****II. MOTIVATION**

Today's monitoring tools still have noticeable weak spots that hold back real progress in the field. They almost never explain how an AI reached its conclusion, so engineers are left guessing and second-guessing every alert. They struggle to separate a faulty sensor that is simply drifting from a deeper system problem that might affect the whole operation. There is no single, easy-to-read health score that tells the full story at a glance, making it hard to prioritize what needs attention first. And visuals are often too basic or cluttered to support fast decisions under pressure when time really matters.

This project was created to close those gaps once and for all. It delivers a simple, affordable, transparent, and scalable AI solution that makes time-series health monitoring more reliable, understandable, and genuinely useful for operators in the field. Instead of treating AI as a mysterious black box, the framework opens it up so everyone can trust the results and act with confidence. The low-cost design also means smaller teams and organizations can benefit without needing expensive hardware or cloud subscriptions. In the end, the motivation is straightforward: build something that actually helps people keep critical systems running smoothly, safely, and with fewer surprises.

III. EXISTING SYSTEM

Present-day solutions generally fall into three familiar categories. Threshold-based systems rely on preset limits to raise alerts; they are easy to set up but cannot adjust when conditions change or when normal patterns slowly drift over time. Black-box machine learning models deliver strong accuracy yet give no clue about why a result appeared, leaving teams to hunt for answers manually. Predictive maintenance tools try to forecast future failures but rarely explain anything in real time, so operators still face uncertainty when something unexpected shows up.

Across all three, the same core problems remain:

- No precise fault localization
- No unified health index
- Poor overall interpretability that limits trust and practical use in critical operations.

IV. PROPOSED SYSTEM

The proposed framework offers a complete, easy-to-use pipeline that turns raw sensor data into clear, actionable health insights for real teams in safety-critical environments.

A. Data Ingestion – Loading Multivariate Time-Series CSV Files The system starts by directly reading standard CSV files containing multiple sensor readings recorded over time, requiring no complex setup or conversion.

B. Preprocessing – Cleaning, Normalizing, and Creating Sliding Windows Raw data is cleaned by filling missing values intelligently, removing noise, normalizing all sensor readings to the same scale, and splitting the continuous series into fixed-size sliding windows that preserve temporal relationships.

C. Feature Engineering – Extracting Statistical and Temporal Features Meaningful features are pulled from each window, including mean, standard deviation, variance, range, and time-based patterns that reveal trends and sudden changes.

D. Anomaly Detection – Using Isolation Forest and Autoencoder Models Two unsupervised models run together: Isolation Forest quickly isolates unusual points, while the autoencoder learns normal patterns and flags deviations through reconstruction error.

E. Fault Diagnosis – Classifying Sensor-Level vs System-Level Anomalies Every anomaly is automatically examined and labelled as either a localized sensor fault or a broader system-wide issue to guide immediate troubleshooting.

F. Health Index Computation – Calculating a Weighted 0–100 Score A single intuitive health index from 0 to 100 is produced by assigning importance weights to each sensor, allowing easy tracking of overall system degradation over time.

G. Explainability Module – Applying SHAP for Transparent Insights SHAP values are generated for every prediction, showing exactly which features contributed most and turning every alert into a clear, trustworthy explanation.

H. Visualization Dashboard – Interactive Streamlit Interface All outputs are displayed in one clean, interactive dashboard with health gauges, trend graphs, anomaly timelines, and SHAP plots for instant understanding.



Fig.1. Architecture of the proposed monitoring system.

V. SYSTEM MODULES

Every part of the system has its own focused job that works smoothly with the others. Data processing cleans incoming records, fills missing numbers without introducing bias, normalizes scales so sensors can be compared fairly, and segments the series so models can work efficiently on manageable chunks. Feature engineering turns raw windows into meaningful summaries—means, variances, trends—that reveal how the system is really behaving beneath the surface. Anomaly detection applies Isolation Forest and autoencoders to flag anything unusual that might indicate early trouble. Fault diagnosis then classifies each flagged event as either a sensor-specific issue that can be fixed locally or a broader system problem that needs wider attention. Health index computation blends all sensor states into a single weighted score from 0 to 100 that updates continuously. Explainability uses SHAP to produce plain-language reasons for every alert, showing which sensors or features pushed the decision. The visualization dashboard puts it all on screen in an interactive format, letting users watch live trends, review past anomalies, zoom into details, and understand root causes without needing extra tools or training.

VI. IMPLEMENTATION

A. Data Preprocessing and Segmentation

The process begins with preprocessing the input time-series data to ensure consistency and quality. Missing values are handled using imputation techniques, and normalization is applied to scale sensor readings into a uniform range.

To capture temporal dependencies, the data is segmented using a sliding window approach. This converts continuous time-series data into fixed-length sequences, enabling the system to analyze patterns over time rather than individual data points. This step is essential for detecting temporal anomalies and improving model performance.

B. Feature Extraction

Following preprocessing, statistical features are extracted from each segmented window to represent system behaviour. These features include measures such as mean and variance, which capture the central tendency and variability of the data.

This transformation reduces noise and dimensional complexity while preserving important characteristics of the system. As a result, the extracted features provide a more informative input for anomaly detection algorithms.

C. Anomaly Detection

The core of the implementation lies in detecting anomalies using machine learning techniques. Isolation Forest is employed to identify abnormal patterns by isolating data points based on random partitioning.

The anomaly score is computed as:

$$s(x,n) = 2^{-\frac{E(h(x))}{c(n)}}$$

where $E(h(x))$ represents the expected path length required to isolate a data point, and $c(n)$ is a normalization factor. Anomalous points tend to have shorter path lengths, resulting in higher anomaly scores. Additionally, an Autoencoder model is used to detect anomalies based on reconstruction error. The model learns normal patterns in the data and flags deviations when reconstruction error exceeds a threshold, enabling robust detection of complex anomalies.

D. Fault Diagnosis

After detecting anomalies, the process proceeds to identify their origin. The system analyses the distribution of anomalies across multiple sensors to determine whether the issue is localized or system-wide.

If anomalies are confined to specific sensors, they are classified as sensor-level faults. In contrast, anomalies affecting multiple sensors simultaneously indicate system-level faults. This step provides meaningful insights into the root cause of abnormal behaviour.

E. Health Index Computation

To quantify overall system condition, a Health Index is computed using a weighted aggregation of sensor states:

$$HI = \sum_{i=1}^n w_i \cdot S_i$$

where S_i represents the health score of each sensor and w_i represents its importance. The final score is scaled between 0 and 100 to provide an intuitive measure of system health.

F. Explainability and Interpretation

The final stage of the process focuses on interpreting model predictions using SHAP. This technique assigns contribution values to each feature, enabling the system to explain why a particular anomaly was detected.

$$f(x) = \phi_0 + \sum_{i=1}^n \phi_i$$

where ϕ_i represents the contribution of each feature. This enhances transparency and allows users to understand the factors influencing system behavior, making the framework suitable for safety-critical applications.

VII. EXPERIMENTAL RESULTS



Fig.2. Anomaly detection output plotted against time using sensor data

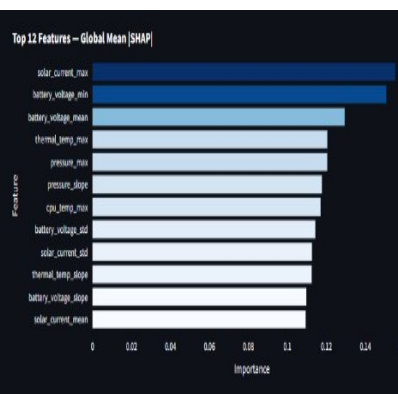


Fig.3. SHAP visualization highlighting the Contribution of each sensor (feature) to anomaly detection



Fig.4. The figure displays the generated system report summarizing key outputs of the framework

Volume 15, Issue 4, April 2026**|DOI: 10.15680/IJRSET.2026.1504227|**

Fig.2. States that the time-series sensor telemetry with anomaly markers overlaid. The red vertical lines represent detected anomaly windows across different sensor channels such as battery voltage, solar current, and CPU temperature. It can be observed that anomalies correspond to sudden deviations, such as voltage drops, current spikes, and temperature rises, indicating abnormal system behaviours over time.

Fig.3. States that the SHAP-based feature importance visualization, highlighting the contribution of each feature to anomaly detection. Features such as *solar_current_max* and *battery_voltage_min* show higher importance values, indicating their strong influence in identifying anomalies. This visualization helps in understanding which sensor parameters contribute most significantly to system faults

Fig.4 states that the generated textual report summarizes key system metrics, including the Health Index, anomaly frequency, and fault severity. The Health Index of 88.7 indicates a nominal but slightly degrading system condition. The report also identifies *battery voltage* as the primary fault source, supported by high SHAP contribution values and Z-score severity.

Together these outputs prove the framework does more than just detect trouble—it explains it in ways that help people make faster, smarter decisions in high-stakes environments. The combination of accurate detection, clear fault classification, and transparent explanations reduces downtime, lowers maintenance costs, and builds real confidence in the AI's recommendations

VIII. CONCLUSION

The explainable AI framework successfully brings together data preparation, feature creation, anomaly detection, fault diagnosis, health scoring, and SHAP explanations into one smooth pipeline for time-series monitoring. It delivers accurate alerts along with easy-to-understand reasons and a practical health score that tracks slow changes over time, giving operators a complete picture they can act on immediately. The modular design makes it straightforward to adapt for different industries or data sources, and the human-friendly outputs lower the barrier for real-world adoption. Looking ahead, the plan is to add real-time streaming support so the system can monitor live data feeds without delay, bring in stronger models like LSTM or Transformers for even better time awareness in long sequences, and include adaptive learning so the framework keeps improving automatically as new data arrives. Further testing on larger real-world datasets from actual industrial sites and wider use in IoT networks, manufacturing plants, and aerospace telemetry are also on the roadmap. These next steps will make the system even more robust and ready for the growing demands of modern safety-critical operations.

IX. ACKNOWLEDGMENT

The authors sincerely thank Mrs. E. Benitha Sowmiya for her steady guidance, insightful feedback, and constant encouragement throughout this research. We are also grateful to the Department of Computer Science and Engineering at Bharath Institute of Science and Technology and Bharath Institute of Higher Education and Research for providing the computational resources, lab facilities, and supportive environment that made this work possible.

REFERENCES

- [1] Artificial Intelligence (XAI) on Time Series Data: A Survey,” *arXiv preprint*, 2021.
- [2] J. Sipple and A. Youssef, “A General Purpose Method for Applying Explainable AI for Anomaly Detection,” *arXiv preprint*, 2022.
- [3] M. Tidriri, N. Chatti, S. Verron, and T. Tiplica, “Bridging Data-Driven and Model-Based Approaches for Process Fault Diagnosis and Health Monitoring: A Review,” *IEEE Access*, vol. 10, 2022.
- [4] L. C. Brito, J. P. S. Catalão, and R. J. Bessa, “Fault Diagnosis Using Explainable Artificial Intelligence: A Transfer Learning Approach,” *Expert Systems with Applications*, vol. 219, 2023.
- [5] E. Brusa, L. Ferretto, and A. Zumpano, “Explainable Artificial Intelligence for Machine Fault Diagnosis,” *Applied Sciences*, vol. 13, no. 4, 2023.
- [6] A. Alharthi, M. Alam, and S. Khan, “Explainable AI for Sensor Signal Interpretation in Health Monitoring Systems: A Review,” *IEEE Access*, vol. 12, 2024.

- [7] J. Cação, J. Santos, and M. Antunes, “Explainable AI for Industrial Fault Diagnosis: A Systematic Review,” *Journal of Industrial Information Integration*, vol. 34, 2025.
- [8] G. Pang, C. Shen, L. Cao, and A. V. D. Hengel, “Deep Learning for Anomaly Detection: A Comprehensive Review,” *ACM Computing Surveys*, vol. 54, no. 2, 2021.
- [9] H. Blázquez-García, A. Conde, U. Mori, and J. A. Lozano, “A Review on Outlier/Anomaly Detection in Time Series Data,” *ACM Computing Surveys*, vol. 54, no. 3, 2021.
- [10] Y. Lei, N. Li, L. Guo, N. Li, T. Yan, and J. Lin, “Machinery Health Prognostics: A Systematic Review from Data Acquisition to RUL Prediction,” *IEEE Transactions on Industrial Electronics*, vol. 69, no. 3, 2022.
- [11] Z. Chen, C. K. Yeo, B. S. Lee, and C. T. Lau, “Autoencoder-Based Anomaly Detection in Time-Series Data: A Survey,” *IEEE Access*, vol. 9, 2021.
- [12] R. Chalapathy and S. Chawla, “Deep Learning for Anomaly Detection: A Survey,” *arXiv preprint*, 2019.



INTERNATIONAL
STANDARD
SERIAL
NUMBER
INDIA



INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN SCIENCE | ENGINEERING | TECHNOLOGY

 9940 572 462  6381 907 438  ijirset@gmail.com



www.ijirset.com

Scan to save the contact details